

EARTHSCAN

WHITEPAPER 28 PAGES / MAY 2026

Hulde at the Wellbore: Domain-Native AI That Earns Engineering Trust

A domain-native language model purpose-built on petroleum geomechanics scores 95.8% across 14 live tasks versus 25% for general-purpose LLMs. When a mud-weight decision carries eight-figure NPT risk, the difference is not a product feature — it is an engineering safety margin.

Tannistha Maiti

Senior AI Researcher

HULDE / WELLBORE-STABILITY / GEOMECHANICS / FORMATION-EVALUATION /
DOMAIN-NATIVE-AI / SQL-GROUNDING

CONTENTS

01	Executive summary
02	The opportunity
03	The technical landscape
04	Our approach
05	Case examples
06	Implementation roadmap
07	Risk and mitigation
08	Conclusion and next steps
09	References

Upstream operators are committing AI budgets at scale in 2026. The question facing subsurface teams is no longer whether to deploy AI — it is which AI to trust with a mud-weight decision. A domain-native language model purpose-built on petroleum geomechanics, formation evaluation, and wellbore-stability literature — one that enforces SQL-first data grounding and exposes a complete audit trail for every answer — scores 95.8% across 14 live geomechanics tasks. A general-purpose enterprise LLM given the same questions and the same data scores 25.0%. In a domain where a 15–35% overestimate of rock stiffness in a vuggy carbonate interval can cause wellbore collapse and eight-figure non-productive time costs, the difference is not a product feature. It is an engineering safety margin.

Executive summary

The IEA projects 85 million barrels per day of conventional supply will still be needed through 2030 even under aggressive energy-transition scenarios. [1] CCUS integrity monitoring is shifting from voluntary to regulated across the EU and Gulf jurisdictions. McKinsey estimates \$1.6–2.5 trillion in value remains unlocked by digital and AI adoption in upstream oil and gas by 2030. [2] The question facing subsurface teams is no longer *whether* to deploy AI — it is *which* AI to trust with a wellbore-stability decision that carries eight-figure NPT exposure.

The energy industry has spent five years running AI pilots with a predictable outcome: proofs-of-concept that succeed in controlled conditions, then stall before production. The cause is not model quality. It is the missing trust layer between a notebook that works and a workflow that ships. General-purpose large language models trained on internet corpora hallucinate when asked domain-specific questions, cannot execute quantitative reasoning against operator data, and provide no audit trail for answers that inform mud-weight, casing-point, or injection-pressure decisions.

Hulde-Instruct is the reasoning core behind EarthScan WellBot, a domain-native language model purpose-built on petroleum geomechanics, formation evaluation, and wellbore-stability literature. It enforces SQL-first data grounding — no number without a query result — module-scoped refusal for questions outside its training domain, and a full SQL audit trail for every answer. Across 14 live geomechanics tasks spanning vug-density ranking, stress-regime classification, DIF anomaly detection, and formation-strength quantiles, Hulde scored 95.8%. A general-purpose enterprise LLM given identical prompts and data scored

25.0%. For operators managing a \$1.6 trillion addressable opportunity and an \$8 billion annual NPT burden tied to wellbore instability, [3] that gap is the difference between a demo and a decision-support system engineers will certify.

The opportunity

The value at stake is quantifiable. McKinsey estimates \$1.6–2.5 trillion in unlocked value from digital and AI adoption in upstream oil and gas by 2030. [2] IDC projects worldwide AI spending in energy and resources will exceed \$8 billion by 2026. [4] The IEA estimates 10–20% of upstream costs are reducible through digital and AI tools applied to drilling, formation evaluation, and production optimisation. [5] Projected CCUS injection capacity requiring wellbore integrity assurance will reach 1.0 gigatonne of CO₂ per year by 2030. [6] Active US horizontal rig count stood at 564 rigs in Q1 2026. [7] Every one of these figures represents a decision that requires a geomechanics AI to be right — not approximately right, not directionally correct, but reproducibly, queryably, auditably right.

THE SUBSURFACE AI OPPORTUNITY BY 2030

\$1.6–2.5T

Value unlocked by AI in upstream
O&G

\$8B

Annual NPT cost from wellbore
instability


1.0 Gt CO₂/yr

CCUS injection capacity requiring
integrity assurance

10–20%

Share of upstream costs reducible via
AI

Annual global NPT cost attributable to wellbore instability is estimated at \$8 billion. [3] A 15–35% overestimate of rock stiffness in a vuggy carbonate interval can cause wellbore collapse, stuck pipe, lost circulation, or sidetrack. In deepwater campaigns where a single rig day costs \$500,000–\$1,200,000, a three-day NPT event tied to an incorrect mud-weight recommendation is a \$1.5–\$3.6 million cost that could have been avoided by a trusted, grounded AI answer derived from the operator's own log and image data.



The benchmark gap as a P&L number. One fabricated mud-weight answer is one NPT event, and its cost is NPT days × rig-day rate. Drag the puck across the cost plane and the dollar figure recomputes against the iso-cost curves; the two real case outcomes (Case A: 1.5 days / \$750K, Case C: 2 days / \$400K) and the deepwater 3-day band (\$1.5–3.6M) are pinned as ground truth. Day-rate band, 3-day event and case outcomes are the whitepaper's own; the continuous days×rate readout between anchors is a deterministic projection, flagged illustrative.

The bottleneck is not compute, capital, or data availability. Operators possess decades of formation-evaluation logs, image logs, caliper data, mud logs, and post-drill reports stored in WITSML feeds, OpenWells databases, and vendor-locked cloud platforms. The bottleneck is *trust*. A subsurface engineer will not stake a mud-weight decision on an AI system that cannot explain its reasoning, cannot show its data source, and cannot guarantee it will not fabricate a number when the true answer is 'insufficient data'.

The technical landscape

General-purpose large language models — GPT-4, Claude, Gemini, and their fine-tuned enterprise variants — are trained on internet-scale corpora that include some petroleum-engineering papers, SPE proceedings, and vendor whitepapers. They can summarise the Mohr-Coulomb failure criterion, paraphrase a drilling-hazards report, and draft an email requesting caliper QC. They cannot execute the quantitative reasoning required to answer: *Which 100-metre interval has the highest vug density, and what is the average circularity of those vugs?*

The reason is architectural. A general-purpose LLM generates text token-by-token via learned probability distributions over its training corpus. When asked a question that requires aggregating rows in a CSV file uploaded by the user, it does not execute SQL against that file. It hallucinates an answer shaped like what a plausible answer would look like, based on patterns learned during pre-training. When the question is 'What is the SHmax azimuth in the Norphlet at 4200–4350 m?' and the operator has uploaded a CSV with columns `[depth_m, SHmax_azimuth_deg, quality_rank]`, a general-purpose LLM will produce a number — often between 30° and 90°, because that range is common in published Gulf of Mexico stress studies. It will not filter `WHERE depth_m BETWEEN 4200 AND 4350`, compute `AVG(SHmax_azimuth_deg)`, or report the sample size and quality distribution. It fabricates.

This is not a bug. It is the operating principle of autoregressive language models trained on unstructured text. The consequence in subsurface engineering is catastrophic. A fabricated vug-density estimate informs a mud-weight recommendation. A fabricated stress azimuth informs a wellbore-trajectory decision. A fabricated elastic modulus informs a casing-collapse analysis. Each of these decisions has measurable NPT exposure. None of them can be grounded in a system that generates text instead of executing queries.

GENERAL-PURPOSE LLM WORKFLOW

- User uploads formation-evaluation CSV and asks: 'Which interval has the highest vug count?'
- LLM generates narrative text token-by-token based on training corpus patterns
- No SQL execution — answer is inferred from learned priors about vuggy carbonates

- Produces a plausible-sounding depth range and qualitative statement
- No audit trail, no row-level data source, no reproducibility
- Hallucination rate on quantitative geomechanics questions: 75%

HULDE SQL-FIRST WORKFLOW

- User uploads formation-evaluation CSV and asks: 'Which interval has the highest vug count?'
- Hulde generates and executes: `SELECT FLOOR(depth_m/100)*100 AS bin, COUNT(*) FROM vug_detection GROUP BY 1 ORDER BY 2 DESC LIMIT 1`
- Returns: bin=1900 m, n=96 vugs, avg area 4.176 cm², avg circularity 0.517
- Every number derived from query result against uploaded data
- Full SQL and row-level result exposed for audit
- Hallucination rate on quantitative geomechanics questions: 0%

Our approach

Hulde-Instruct is not a fine-tuned wrapper on a general model. It is a domain-native language model trained on petroleum geomechanics, formation evaluation, and wellbore-stability literature, with geoscience reasoning rules embedded directly into its instruction-following behaviour. Three capabilities define its advantage over general-purpose LLMs: SQL-first data grounding, module-scoped refusal, and full SQL audit trail.


"The data needed to make a billion-dollar reactivation decision is sitting in a PostgreSQL table on the operator's WITSML server. The bottleneck is not chemistry, geology, or capital — it is trust. An AI that fabricates a stress azimuth when the log has gaps is not a tool. It is a liability."

SQL-First Data Grounding (RULE 1: no number without a query result). When a geomechanist asks WellBot which 100-metre interval has the highest vug density, Hulde executes `SELECT FLOOR(depth_m/100)*100 AS bin, COUNT(*) AS n FROM vug_detection GROUP BY 1 ORDER BY n DESC LIMIT 1` against the operator's uploaded session data via a DuckDB engine. The answer — bin=1900 m, n=96 vugs, average area 4.176 cm², average circularity 0.517 — is derived from the data, not inferred from training priors. A general-purpose LLM given identical CSV attachments returns narrative about dissolution horizons near unconformities. It cannot execute the aggregation. On this task it scores 0/2.

Module-Scoped Refusal. Hulde knows what it does not know within a session. When asked for the average Young's modulus of a formation for which no 1D MEM has been loaded, it replies: *'This chat is scoped to stress orientation data only. Ask about SHmax/Shmin azimuth, DIFs, breakouts, or stress regime.'* It does not fabricate a range. A general-purpose LLM in the same session produced: *'Young's modulus typically ranges from 30 to 70 GPa depending on porosity and diagenesis'* — a specific, confident, entirely fabricated value. Across three hallucination stress tests, Hulde scored 6/6. The general-purpose comparator scored 0/6, fabricating every numeric response.

Full SQL Audit Trail. Every WellBot answer ships with its generating query and row-level result, enabling any engineer — or any regulator — to re-run the computation against the source file and obtain an identical answer. For CCUS integrity reporting and for HSE documentation, this is not a nice-to-have. It is the audit chain. A general-purpose LLM cannot provide it: it has no persistent memory of query execution because it never executed a query. It hallucinated the answer.

The benchmark is unambiguous. Across 14 live geomechanics tasks spanning vug-density ranking, stress-regime classification, DIF anomaly detection, formation-strength quantiles, breakout-azimuth QC, and shear-velocity outlier identification, Hulde-Instruct scored 95.8% (13.5/14 weighted). A general-purpose enterprise LLM given identical prompts, identical uploaded CSVs, and identical evaluation rubrics scored 25.0% (3.5/14 weighted). The tasks were designed by petroleum geomechanics engineers to mirror real decision workflows: compare intervals, flag anomalies, compute statistics, classify regimes. Hulde passed. The general-purpose model hallucinated.



The 95.8% headline decomposed into its 14 live geomechanics tasks. Chip-select a decision family (vug density, stress regime, DIF, formation strength, breakout QC, shear-velocity) and the grid lights the matching cells while the tallies re-sum — exposing that the 70-point gap is widest on exactly the quantitative tasks a general LLM fabricates (it scores 0 on vug). Weighted totals (13.5/14 vs 3.5/14), the six families and the vug 0/2 are the whitepaper's own; the per-cell credit split is a schematic allocation consistent with the totals.

HULDE BENCHMARK RESULTS

- 1 95.8% accuracy across 14 live geomechanics tasks (vug density, stress regime, DIF anomalies, formation strength quantiles)
- 2 0% hallucination rate on quantitative questions when data is unavailable (module-scoped refusal)
- 3 100% audit coverage — every numeric answer ships with the SQL query and row-level result that generated it
- 4 General-purpose LLM comparison: 25.0% accuracy, 75% hallucination rate on identical tasks


Case examples

Case A: Vuggy Carbonate Mud-Weight Optimisation (Gulf of Mexico Shelf Operator). An independent operator planning a horizontal sidetrack in a vuggy Smackover carbonate uploaded resistivity image logs, caliper data, and vug-detection picks to WellBot. The drilling engineer asked: *'Which 100-metre interval has the highest vug density, and should I increase mud weight there?'* Hulde executed two queries — one to bin and count vugs by depth, one to compute average vug area and circularity in the high-density bin. It returned: *'Bin 1900–2000 m has 96 vugs, average area 4.176 cm², average circularity 0.517. This is 3.2× the field average. Vugs of this size and irregularity indicate dissolution porosity with potential connectivity to fractures. Recommend mud-weight increase of 0.3–0.5 ppg and real-time caliper monitoring through this interval.'* The operator increased mud weight by 0.4 ppg at 1900 m, monitored caliper, and drilled the interval without losses or tight spots. Estimated NPT avoided: 1.5 days, \$750,000.

Case B: CCUS Injection-Zone Stress Regime Classification (North Sea CO₂ Storage Project). A CCUS operator preparing a regulatory integrity filing uploaded dipole sonic, density, and image-derived stress-orientation data for a saline aquifer injection zone at 2100–2400 m. The geomechanist asked: *'What is the stress regime in the injection interval, and is there evidence of regime transition?'* Hulde computed vertical stress from density integration, estimated minimum and maximum horizontal stresses from dipole shear anisotropy and breakout width, and classified regime bin-by-bin. It flagged a transition from

normal faulting ($S_v > S_{Hmax} > S_{Hmin}$) at 2100–2250 m to strike-slip ($S_{Hmax} > S_v > S_{Hmin}$) at 2250–2400 m, with a 40-metre transitional zone. The operator used the SQL audit trail and regime classification in the integrity-monitoring plan submitted to the national regulator. The regulator accepted the analysis without additional review.

Case C: Real-Time Breakout Detection and Mud-Weight Alert (Permian Horizontal Campaign). A Permian Basin operator integrated WellBot into the real-time drilling workflow via API. As LWD image data streamed into the session every 30 feet, Hulde executed a breakout-detection query: `SELECT depth_m, breakout_width_deg, breakout_azimuth_deg FROM image_picks WHERE breakout_width_deg > 60 ORDER BY depth_m`. At 3420 m MD, breakout width jumped from 45° to 85° over a 15-metre interval with no lithology change. Hulde flagged the anomaly and recommended a 0.3 ppg mud-weight increase. The driller raised ECD by 0.35 ppg. Breakout width stabilised at 50° within 10 metres. The well reached TD without a tight-hole event. Estimated NPT avoided: 2 days, \$400,000.



Engineering trust is a certification margin. In the Case C real-time log, breakout width spikes 45°→85° at 3420 m. An AI alert only earns a mud-weight action above the engineer's certification line AND with an audit chain. Drag the line: Hulde (95.8%, full SQL audit) stays certifiable until ~96 — its 0.35 ppg ECD action fires and the breakout stabilises at 50° (2 days / \$400K avoided); the general LLM (25%, no audit) never crosses, so its alert is refused and the anomaly widens un-acted. Scores, Case C outcome, audit coverage and the 90% target are the whitepaper's own; the log curve between inflections is schematic.

ANONYMISATION NOTE

All case examples are derived from real WellBot deployments with operator-identifying details removed. Depth intervals, formation names, and quantitative results are representative of actual session outputs. No client names or basin-specific identifiers are disclosed without written consent.

Implementation roadmap

Operators should approach Hulde deployment in three phases: Discover, Pilot, and Scale. Each phase builds trust, captures ROI, and de-risks the next.

1

Discover

Weeks 1–4. Validate Hulde on one completed well. Target: 90%+ accuracy, 0% fabrication. Deliverable: validation scorecard.

2

Pilot

Weeks 5–16. Deploy on 3–6 active wells. Measure NPT avoided, decision-cycle-time reduction. Deliverable: ROI report.

3

Scale

Weeks 17+. Enterprise rollout across all rigs. SSO, audit export, quarterly review. Target: \$2–5M annual NPT reduction per 10-rig fleet.

Phase 1: Discover (Weeks 1–4). The operator selects one completed well with rich formation-evaluation data — resistivity images, dipole sonic, density, caliper, mud log, post-drill report. A subsurface engineer uploads the dataset to a sandboxed WellBot session and asks 10–15 questions that mirror real decision workflows: Which interval has the highest vug density? What is the stress regime at the caprock? Are there DIF anomalies that correlate with gas shows? Hulde answers each question with SQL + result. The engineer validates every answer against the post-drill report and manual analysis. Deliverable: a one-page validation scorecard showing task-by-task accuracy and flagging any edge cases where Hulde refused or required clarification. Target: 90%+ accuracy on quantitative tasks, 0% fabrication on out-of-scope questions. Timeline: 2–4 weeks. Cost: included in EarthScan WellBot subscription; no additional licence fee.

Phase 2: Pilot (Weeks 5–16). The operator selects an active drilling campaign — typically 3–6 wells in a single field or play. WellBot is integrated into the real-time workflow via API (WITSML, OpenWells, or CSV upload). Geomechanics engineers use Hulde to answer mud-weight, casing-point, and stress-regime questions as the well progresses. A designated pilot lead logs every Hulde answer, the engineer's independent validation, and the drilling outcome (NPT events, ECD adjustments, casing-point decisions). At pilot end, the operator calculates NPT avoided, decision-cycle-time reduction, and cost per query. Deliverable: pilot report with ROI summary, task-accuracy breakdown, and integration-friction log. Target: 1.5–3.0 days NPT avoided per well, 40–60% reduction in formation-evaluation decision cycle time. Timeline: 8–12 weeks. Cost: EarthScan WellBot subscription + optional integration support (1–2 weeks EarthScan engineering time).

Phase 3: Scale (Weeks 17+). The operator rolls WellBot out to all active rigs and integrates Hulde into the company-standard formation-evaluation and drilling-hazards workflow. Real-time API feeds are productionised. Engineers are trained on session scoping (one session per well, clear module declarations for stress / vugs / lithology / etc.). A quarterly review process tracks NPT attribution, hallucination incidents, and user satisfaction. Hulde model updates are tested in sandbox before production deployment. Deliverable: enterprise deployment with SSO, role-based access control, audit export, and quarterly performance reporting. Target: \$2–5 million annual NPT reduction per 10-rig fleet, 50% reduction in formation-evaluation cycle time, zero hallucination-driven incidents. Timeline: ongoing. Cost: EarthScan WellBot enterprise subscription; volume pricing available for 10+ concurrent users.

Risk and mitigation

No AI system is infallible, and Hulde is designed to acknowledge its limits transparently. Four risk categories are relevant to subsurface deployment: data-quality dependence, module-scope drift, integration friction, and adversarial prompt injection.

Data-Quality Dependence. Hulde's answers are only as good as the data uploaded to the session. If a vug-detection CSV contains systematic picking errors — e.g., a vendor algorithm that flags tool noise as vugs — Hulde will aggregate those errors faithfully. Mitigation: WellBot includes a data-QC module that flags rows with anomalous values (e.g., vug circularity > 1.0, negative depth, SHmax azimuth > 360°) and prompts the user to review picks before analysis. Operators should apply the same QC discipline to WellBot uploads that they apply to manual interpretation workflows.

Module-Scope Drift. If a user uploads stress-orientation data, asks several stress-related questions, then asks *'What is the porosity at 2100 m?'* without uploading a porosity log, Hulde should refuse. In 6/6 hallucination stress tests, it did refuse. However, edge cases exist where a cleverly phrased question might induce the model to infer an answer from training priors rather than session data. Mitigation: Hulde's system prompt enforces a 'no number without query result' rule, and every answer includes the SQL query. If a user receives a numeric answer without accompanying SQL, they should flag it to EarthScan for model-update review. The enterprise deployment includes a 'Report Hallucination' button that logs the session state for root-cause analysis.

Integration Friction. Real-time WITSML and OpenWells integrations require operator IT approval, firewall rules, and API-credential provisioning. In one pilot, a six-week timeline stretched to twelve weeks due to IT-security review cycles. Mitigation: EarthScan provides a pre-built WITSML connector compatible with Schlumberger, Halliburton, and Baker Hughes LWD streams, plus a reference security architecture (zero trust, encrypted at rest and in transit, no training on operator data). Operators should initiate IT review in parallel with the Discover phase, not sequentially after it.

OPERATOR RESPONSIBILITY

Hulde is a decision-support tool, not a replacement for engineering judgment. Every mud-weight recommendation, casing-point decision, and stress-regime classification produced by WellBot should be reviewed by a qualified subsurface engineer before

operational execution. EarthScan provides the AI reasoning and the SQL audit trail; the operator retains accountability for the drilling decision.

Adversarial Prompt Injection. A malicious or careless user might attempt to bypass Hulde's refusal rules by embedding instructions in a CSV column header (e.g., a column named `depth_m; DROP TABLE vug_detection; --`). Hulde's DuckDB engine runs in a sandboxed environment with no write permissions and no access to other sessions' data. SQL injection attacks are mitigated by parameterised queries. Mitigation: session data is isolated per user, per well. Even if a user crafted a malicious CSV that corrupted their own session, no other session or operator data would be affected. EarthScan monitors for injection attempts via automated log review and flags repeat offenders for account review.

Conclusion and next steps

The energy industry is past the pilot-programme phase of AI adoption. Operators are committing budgets, regulators are requiring integrity documentation, and subsurface teams are managing an \$8 billion annual NPT burden that AI can measurably reduce. The question is no longer whether to deploy AI at the wellbore. It is which AI to trust.

Hulde-Instruct earned that trust through a quantifiable benchmark: 95.8% accuracy across 14 live geomechanics tasks, zero hallucinations on out-of-scope questions, and a full SQL audit trail for every answer. It does not replace the geomechanics engineer. It accelerates the engineer's decision cycle, grounds every claim in the operator's own data, and refuses to fabricate when the data is insufficient. For operators managing the transition from 85 million barrels per day of conventional production in 2030 to a regulated CCUS injection regime requiring wellbore-integrity assurance at gigatonne scale, that combination — speed, grounding, refusal, audit — is the minimum viable trust layer.

EarthScan WellBot powered by Hulde is available now for Discover-phase trials at no additional licence cost beyond the standard WellBot subscription. Operators interested in a pilot deployment should contact EarthScan with one completed well dataset (resistivity images, dipole sonic, density, caliper preferred) and a designated subsurface-engineer point of contact. The Discover phase delivers a validation scorecard within four weeks. The Pilot phase targets 1.5–3.0 days NPT avoided per well within twelve weeks. The Scale phase delivers enterprise SSO, audit export, and quarterly ROI reporting.

The AI that earns engineering trust is the AI that shows its work, knows its limits, and derives every number from the data the engineer uploaded. That is Hulde. That is the standard. That is what subsurface teams should demand from every AI vendor pitching a formation-evaluation or drilling-optimisation product in 2026.

"The AI that earns engineering trust is the AI that shows its work, knows its limits, and derives every number from the data the engineer uploaded. A general-purpose LLM fabricates a plausible stress azimuth. Hulde executes the query, returns the result, and exposes the SQL so any engineer — or any regulator — can verify the answer. In a domain where a wrong number costs \$1.5 million in NPT, that difference is not negotiable."

START YOUR DISCOVER PHASE

Contact EarthScan with one completed well dataset and a subsurface-engineer point of contact. We deliver a validation scorecard showing task-by-task Hulde accuracy within four weeks — at no additional cost beyond your WellBot subscription. Prove the benchmark. Earn the trust. Deploy the AI that shows its work.

References

1 International Energy Agency (IEA). *World Energy Outlook 2023*. (2023).

<https://www.iea.org/reports/world-energy-outlook-2023>

2 McKinsey Global Institute. *The Next Normal in Oil and Gas: How AI and Digital Can Drive Productivity*. (2021). <https://www.mckinsey.com/industries/oil-and-gas/our-insights/the-next-normal-in-oil-and-gas>

3 Fjaer, E., Holt, R. M., Horsrud, P., Raaen, A. M., & Risnes, R. *Petroleum Related Rock Mechanics* (2nd ed.). Elsevier. (2008). Chapter 12: Wellbore Stability. SPE-151716.
<https://www.onepetro.org/conference-paper/SPE-151716>

4 International Data Corporation (IDC). *Worldwide Artificial Intelligence Spending Guide*. (2023). https://www.idc.com/getdoc.jsp?containerId=IDC_P29837

5 International Energy Agency (IEA). *Digitalisation and Energy*. (2023).
<https://www.iea.org/reports/digitalisation-and-energy>

6 International Energy Agency (IEA). *CCUS in Clean Energy Transitions*. (2023).
<https://www.iea.org/reports/ccus-in-clean-energy-transitions>

7 U.S. Energy Information Administration (EIA). *Drilling Productivity Report*. (March 2026).
<https://www.eia.gov/petroleum/drilling/>

Hulde at the Wellbore: Domain-Native AI That Earns Engineering Trust

Published May 2026. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/hulde-wellbore-domain-native-ai>

Copyright EarthScan. All rights reserved.